



**Development of Cislunar Space Logistics
Networks for Satellite Constellation Support
Using Event-Driven Generalized
Multi-Commodity Network Flows**

THESIS

Alexander R. Collins, 1st Lieutenant, USAF
AFIT-ENY-MS-19-D-048

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION IS UNLIMITED

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government.

AFIT-ENY-MS-19-D-048

Development of Cislunar Space Logistics Networks for Satellite Constellation
Support Using Event-Driven Generalized Multi-Commodity Network Flows

THESIS

Presented to the Faculty

Department of Aeronautics and Astronautics

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Astronautical Engineering

Alexander R. Collins

1st Lieutenant, USAF

November 2019

DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION IS UNLIMITED

AFIT-ENY-MS-19-D-048

Development of Cislunar Space Logistics Networks for Satellite Constellation
Support Using Event-Driven Generalized Multi-Commodity Network Flows

THESIS

Alexander R. Collins
1st Lieutenant, USAF

Committee Membership:

Maj Joshua A. Hess
Chair

Lt Col Bruce A. Cox
Member

Lt Col Kirk W. Johnson
Member

Ms. Kerianne Hobbs
Member

Abstract

As space becomes an increasingly congested domain, the risk of damage to satellite constellations is increasing. In response, there is an increasing need for capabilities for unmanned repair, refueling, and reconstitution (R3) of those constellations. Cislunar orbits offer a promising storage and low-cost transfer solution for on-orbit service vehicles and replacement satellites to leverage those capabilities. This research makes use of mixed-integer linear programming-based logistics models to determine the situations in which a cislunar mission architecture would offer a cost-effective alternative to Earth-based R3.

The network models presented in this research make use of the latest developments in Event-Driven Generalized Multi-Commodity Network Flows (ED-GMCNF), a new method of optimization that enables variable time steps between events. This research combines a new version of an ED-GMCNF with cislunar trajectory optimization to evaluate both the feasibility of cislunar orbits as well as the potential effects of lunar fuel production on R3 costs.

This investigation finds, through an exhaustive numerical simulation campaign, that cislunar logistics networks provide cost-effective means of R3 regiments for Earth-orbiting and cislunar satellites when a lunar fuel supply is taken into consideration. The ED-GMCNF methodology also offers a promising foundation for future work in the mission planning field.

Acknowledgements

I would like to thank my parents and sister for keeping me motivated. I would like to thank my first advisor, Lt Col Kirk Johnson, for his support and direction. I would like to thank my second advisor, Maj Hess, for being willing to step in and help me across the finish line. I would like to thank the other members of my committee for taking the time to provide feedback. I would also like to thank the incredibly supportive folks at AFRL's aerospace systems directorate who allowed me the time and flexibility to get a degree. Finally, I would like to thank the office coffee machine for keeping me awake.

Alexander R. Collins

Table of Contents

	Page
Abstract	iv
Acknowledgements	v
List of Figures	viii
List of Tables	x
 I. Introduction	 1
1.1 Motivation	1
1.2 Overview	3
1.3 Problem Statement	3
1.4 Thesis Overview	4
1.5 Contributions of this Research	5
 II. Background	 6
2.1 The Circular Restricted 3-Body Problem	6
2.1.1 The N-Body Problem	6
2.1.2 Approximate Solution of the CR3BP	8
2.2 Lagrange Points and Stability	11
2.3 Orbits of Interest	14
2.4 Generalized Multi-Commodity Network Flows	17
2.5 Mixed Integer Linear Programming	22
2.6 Sparse Non-Linear Trajectory Optimization	24
2.7 Literature Review	25
2.8 Contributions of this Research	27
2.9 Summary	27
 III. Methodology	 28
3.1 Assumptions and Equations	28
3.1.1 Propellant Mass Fraction	28
3.1.2 Inert Mass Fraction	29
3.2 Transformation Matrix	30
3.3 Proof of Concept	31
3.4 Test Problem Results	34
3.5 Main Network	39
3.5.1 Nodes and Arcs	40
3.5.2 Delta-V and Time of Flight Information	40
3.6 Main Network Parameters	43
3.7 Methodology Summary	48

	Page
IV. Results and Analysis	49
4.1 Earth-Based Refueling Only	50
4.2 Lunar ISRU Included	54
4.3 L1 Lyapunov Constellation Included	59
4.4 Miscellaneous Parameter Variations	61
4.5 Final Analysis of Results	69
4.6 Chapter Summary	69
V. Conclusions and Future Work	70
5.1 Conclusions	70
5.2 Limitations of the ED-GMCNF	70
5.3 Future Work	75
5.4 Contributions of this Research	78
5.5 Summary	78
A. Main Network delta-V and TOF Table	79

List of Figures

Figure		Page
1	CR3BP in Synodic Rotating Coordinate Frame	10
2	Earth-Moon Lagrange Points	13
3	Zero Velocity Curves	16
4	TE-GMCNF and ED-GMCNF Examples	18
5	Linear Programming Cost Function	23
6	Test Network	33
7	Test Problem 1st Event Layer	35
8	Test Problem 2nd Event Layer	35
9	Test Problem 3rd Event Layer	36
10	Test Problem 4th Event Layer	36
11	Test Network with “Shortcut” Arc	38
12	L2 Lyapunov to LLO Transfer	42
13	LLO to GEO Transfer	42
14	Example Pareto Front	44
15	Earth Fuel Only Trial 1	52
16	Earth Fuel Only Trial 2	52
17	Earth Fuel Only Trial 3	53
18	Earth Fuel Only Trial 4	53
19	Lunar ISRU Trial 1	55
20	Lunar ISRU Trial 2	55
21	Lunar ISRU Trial 3	57
22	Lunar ISRU Trial 4	57

Figure		Page
23	Lunar ISRU Trial 5	58
24	Lunar ISRU Trial 6	58
25	L1 Constellation Trial 1	60
26	L1 Constellation Trial 2	60
27	L1 Constellation Trial 3	62
28	L1 Constellation Trial 4	62
29	L1 Constellation Trial 5	63
30	Integer Feasibility Modification Trial 1.....	64
31	MIPGap Modification Trial 1.....	65
32	MIPGap Modification Trial 2.....	65
33	MIPGap Modification Trial 3.....	66
34	MIPGap Modification Trial 4.....	66
35	MIPGap Modification Trial 5.....	67
36	Satellite Mass Modification Trial 1	68
37	Satellite Mass Modification Trial 2	68
38	Network Solution with “Repair Gap”	73
39	Flow Generation Loop Error	75
40	Extra Satellites Removed Example 1	77
41	Extra Satellites Removed Example 2	77

List of Tables

Table		Page
1	Test Network Nodes	31
2	Main Network Nodes	46
3	Main Network Commodities	46
4	Main Network Demands	50

I. Introduction

1.1 Motivation

The Department of Defense (DoD) overwhelmingly recognizes the nature of space as an increasingly contested environment. The unclassified summary of the 2018 National Defense Strategy explicitly outlines the fact that space has become both critical to military and commercial activity, and increasingly threatened by US adversaries [33]. Publications by both the Defense Intelligence Agency (DIA) and the National Air and Space Intelligence Agency (NASIC) go into greater detail about the nature of threats to space activity [9], [6]. These threats take the form of a full spectrum of attacks ranging from reversible, relatively low-cost tactics such as ground station jamming, to catastrophic, irreversible attacks such as kinetic kill vehicles and nuclear detonations [9]. Space is also becoming an increasingly crowded environment as new actors gain access, especially with small satellites. There are currently hundreds of CubeSats in Earth orbits, and roughly one in five have failed to de-orbit, posing a hazard [8]. Cislunar space, in contrast, is almost untouched, with an upcoming mission from Cornell University being one of the first uses [14]. This space is unique from Earth-orbiting space in that gravitational forces from both the Earth and Moon affect a spacecraft's orbit. As the volume of both active satellites and debris in orbit grows, the threat of an intentional impact grows with it. Several methods are being tested for "cleaning up" space, including a harpoon satellite successfully deployed in

February of 2019 as part of a British-led effort [2]. However, missions such as this are still very much experimental ventures, and with an estimated 8,000+ metric tons of debris in orbit [2], satellites in commonly-used orbits such as GEO will experience increasing risk of collision for the foreseeable future. This will mean more frequent replacements are required for constellations, and that such replacements, if stored on-orbit, would be vulnerable to damage themselves. A cislunar network for low-risk storage and rapid replenishment is a potential solution. Cislunar orbits would also have the potential to allow for lower-delta V inclination changes [4], which could reduce the servicing cost of high-inclination satellites and constellations. On-orbit Repair, Refueling, and Reconstitution (R3) capabilities also present a substantial opportunity for cost reduction. Even with the advent of reusable launch vehicles, the cost of launching material to orbit remains high, at an estimated 8,000 USD per kilogram for a GEO launch using the Falcon 9 [39]. Combined with the inherently semi-random nature of replenishment (lifetimes can be predicted [13], but accidents still happen [5]), launch costs can lead to very high maintenance costs, as mission planners must pay a premium for short-notice launches [23]. In this regard, there are significant cost benefits to having spare satellites launched from Earth in bulk at predictable times.

As of 2019, current proposed strategies for on-orbit R3 include both in-plane (in the plane of the constellation being serviced) and out-of-plane spare satellites acting simultaneously [23]. However, these strategies are also heavily focused on Earth-based resupply, meaning that all spare parts, satellites, and propellant must be launched from Earth. The incorporation of cislunar orbits into a logistics network has also been explored, but efforts tend to focus on the Moon as a stepping stone to Mars or in exploration of the Moon itself [22], [17]. Thus far, the concept of a military space logistics network has not been explored in great detail. The optimization of cislunar

orbits, however, has been explored in detail by Dahlke, Brick, and Ostman [12] [3] [35], and elements of this optimization have been incorporated here. One of the most complete accounts of current on-orbit servicing efforts is the April 2019 list by the Aerospace Corporation [24], which includes missions from both the public and private sectors, such as the National Aeronautics and Space Administration’s (NASA’s) Restore-L and Northrop-Grumman’s Mission Extension Vehicle. Unfortunately, save for NASA’s Lunar Gateway, none of the systems discussed in that document have been designed for extended operations in cislunar space.

A cislunar logistics network for constellation replenishment should be investigated because it may offer substantial benefits in both cost and performance to current replenishment methods.

1.2 Overview

This investigation attempts to determine optimal logistics networks that could be used to replenish and resupply satellite constellations. Multiple constellations have been explored, as well as the possibility of in-situ resource utilization (ISRU) to lower costs through the manufacture of hydrogen-based propellant [11]. The focus of this research is on cislunar space. The inclusion of cislunar trajectories allows for the exploration of several potentially useful orbits that could provide distinct advantages over 2-body, Earth-centric orbits. If implemented, a cislunar logistics network could offer not only potential cost savings, but a significant increase in survivability.

1.3 Problem Statement

This investigation attempts to answer two central research questions:

- Can Event-Driven Multicommodity Network Flows (ED-GMCNFs) be used to optimize R3 mission architectures?

- Under what circumstances does the use of cislunar orbits for R3 become cost-effective? Possible scenarios in this investigation are the R3 of constellations with large inclination differences, the use of hypothetical lunar-derived propellant, and the presence of cislunar constellations.

An optimal solution to a network is defined in this research as the set of transfer orbits which minimize the propellant cost of providing replacement satellites and propellant to a constellation while remaining within maximum time of flight requirements. These scenarios have been modeled using the Gurobi software suite [16] and mixed-integer linear programming (MILP) models. Finding these optimal solutions may enable mission planners to take advantage of these alternative orbits.

1.4 Thesis Overview

This investigation attempts to find the optimal solution to several cislunar logistics networks. The use of cislunar space will enable the use of trajectories that can offer significant military advantage to satellites. By developing a method of generating these solutions, future mission planners can quickly determine the viability of both satellite constellations and replenishment and resupply strategies. This thesis presents the investigation as follows:

- Chapter 1 introduces the problem and motivation for solving it.
- Chapter 2 covers the mathematical background of the investigation. Specifically, chapter 2 details the formulation and use of the circular restricted three-body problem (CR3BP), and the use of MILP as they apply to space logistics networks. In addition to the mathematical background, a brief literature review is provided, detailing some of the previous work in these fields.

- Chapter 3 describes the methodology of the investigation, including the initial development of the solver, the formulation of the logistics networks to be tested, and the methodology of the tests themselves. Chapter 3 also presents the results of a test network and the limitations revealed.
- Chapter 4 presents and analyzes the results of primary trials.
- Chapter 5 summarizes the conclusions that can be drawn from the results. Chapter 5 also presents the key limitations of the ED-GMCNF and how those might be removed or mitigated in future work.

1.5 Contributions of this Research

This research represents the first formulation of a space-centric ED-GMCNF with the capacity to select from fuel and time-optimized arcs between the same node pair. In this work, a significant reduction in independent variables has been made compared to the previous work of Ho [21]. This research is also the first of these ED-GMCNF networks to model a reusable launch vehicle.

II. Background

This chapter presents the technical background of the four primary aspects of this investigation: the circular restricted 3-body problem (CR3BP), the modeling of space logistics networks, the use of discrete linear programming as an optimization tool, and the optimization of cislunar trajectories. For each of these areas, this chapter will cover the mathematical basis of the method including relevant historical data, and a discussion of how it's applied to the investigation. This chapter concludes with a discussion of previous literature related to the use of MILP-based logistics networks.

2.1 The Circular Restricted 3-Body Problem

This section gives an overview of the Circular Restricted 3-Body Problem (CR3BP). The CR3BP is a mathematical method for modeling orbits in which 2 objects simultaneously affect the motion of a third object of comparatively negligible mass, such as when a satellite is moving in proximity to both the Earth and the Moon. Historically, the CR3BP has proven difficult to analyze compared to the two-body problem due to the lack of an exact solution, but when simplified it can be used to describe orbits impossible within the solution space of the latter.

2.1.1 The N-Body Problem.

The N-body problem is a method of describing the effect of gravity on the relative motion of objects in orbit. In the N-body problem, each object can be represented by a 6-element state vector describing its position and velocity in 3 spatial dimensions, which means that any N-body problem will have $6N$ unknowns and $3N$ second order equations of motion describing an object's acceleration in the 3 spatial dimensions [44]. The equations of motion for the problem are rooted in the combination of two

known formulae. The first is Newton's second law of motion, which, in an inertial reference frame, is

$$\mathbf{F} = m\mathbf{a} \quad (2.1)$$

where m is an object's mass, $\mathbf{a}$ is its acceleration vector, and $\mathbf{F}$ is the sum of external forces applied to the object. The second is the equation for the force due to gravitational attraction

$$\mathbf{F} = G \frac{m_1 m_2}{r^3} \mathbf{r} \quad (2.2)$$

where G is the universal gravitational constant, m_1 and m_2 are the masses of the bodies attracting each other, and $\mathbf{r}$ is the vector between the bodies' centers of mass, with r being its magnitude (distance between bodies). For N bodies, substituting Newton's equation into the gravitational equation produces:

$$m_i \ddot{\mathbf{r}}_i = \sum_{j \neq i}^N \frac{G m_i m_j}{r_{ij}^3} (\mathbf{r}_j - \mathbf{r}_i) \quad (2.3)$$

where N is the number of bodies under consideration, r_{ij} is the distance between bodies i and j , and j goes from 1 to N , skipping over i , as a body cannot influence itself [44]

Per Newton's third law, the forces between any set of bodies will be equal and opposite, so

$$M \ddot{\mathbf{r}}_c = 0$$

where M is the total mass of all bodies in the system, and $\ddot{\mathbf{r}}_c$ the the second derivative of location vector of the barycenter (combined center of mass). Because the force of gravity can only act on the line between any two objects, angular momentum is also conserved [44]. The N -body problem also assumes that the energy of the system is conserved [44]. These simplifications mean that in a 2-body problem ($N = 2$), an

exact solution is possible; all six state variables $x, y, z, \dot{x}, \dot{y}, \dot{z}$ (the x, y , and z positions and velocities) can be solved for at any point in time. Unfortunately, the solution space becomes more difficult to navigate with the introduction of a third body.

2.1.2 Approximate Solution of the CR3BP.

The CR3BP as used in this investigation consists of two massive primary bodies (the Earth and Moon), which simultaneously influence the motion of a third body of negligible mass (the satellite, tug vehicle, or both, depending on the transportation arc being considered). Since the works of Henri Poincaré in the late 1800's, the CR3BP is known to be unsolvable [44]. When modeling the CR3BP, convention dictates that the units for mass, distance, and time be defined by the properties of the two primary objects. This nondimensionalization greatly simplifies later mathematics. Specifically:

- The mass unit is defined as the sum of the primary masses $m_1 + m_2$.
- The length unit is defined as the distance between the two primary bodies centers of mass a_{12} .
- The time unit is defined as the period of the primary masses orbits around each other $T_{12} = 2\pi\sqrt{\frac{a_{12}^3}{G(m_1+m_2)}} = 2\pi$.

The benefits of this are:

- By defining the smaller mass $m_2 = \mu$, then the mass of the larger body $m_1 = 1 - \mu$.
- Since μ is non-dimensional, it can also be used to describe the distances to the barycenter from the larger mass m_1 (μ) and the smaller mass m_2 ($1 - \mu$).

- Setting $T_{12} = 2\pi$ enables the gravitational constant G to equal 1, per Kepler's 3rd law $T_{12} = 2\pi[\frac{a_{12}^3}{G(m_1+m_2)}]^{1/2}$.

The other assumption made is that this problem is modeled in a rotating coordinate frame fixed to the two primary bodies, with the $\mathbf{s}_1$ axis along the line between them, and the $\mathbf{s}_2$ and $\mathbf{s}_3$ axes at right angles ($\mathbf{s}_3$ points out of the page). The system is diagrammed in Figure 1. [44].

As with the distance, length, and time units, the angular velocity ω forms the default unit of angular velocity, equal to $1\mathbf{s}_3$.

From this point, the equations of motion for the satellite (the body of negligible mass) can be built, starting with defining its position vector in the rotating frame $\mathbf{s}$ as

$$\mathbf{r} = x\mathbf{s}_1 + y\mathbf{s}_2 + z\mathbf{s}_3 \quad (2.4)$$

which means that the velocity vector is

$$\dot{\mathbf{r}} = \dot{x}\mathbf{s}_1 + \dot{y}\mathbf{s}_2 + \dot{z}\mathbf{s}_3 \quad (2.5)$$

and the acceleration vector is

$$\ddot{\mathbf{r}} = \ddot{x}\mathbf{s}_1 + \ddot{y}\mathbf{s}_2 + \ddot{z}\mathbf{s}_3 \quad (2.6)$$

From this, the acceleration of the satellite in the inertial frame, which is necessary in order to use Newton's second law per the N-body problem description above, can be determined as

$${}^i\ddot{\mathbf{r}} = {}^s\ddot{\mathbf{r}} + 2\omega \times {}^s\dot{\mathbf{r}} + \omega \times (\omega \times {}^s\mathbf{r}) \quad (2.7)$$




Figure 1. CR3BP in Synodic Rotating Coordinate Frame: the synodic rotating coordinate frame s is shown, centered on the system's barycenter. [Image credit: Wiesel [44]]

where ${}^i\ddot{\mathbf{r}}$ is the acceleration in the inertial frame and ${}^s\ddot{\mathbf{r}}$ is the acceleration in the rotating frame. These equations simplify to

$${}^i\ddot{\mathbf{r}} = (\ddot{x} - 2\dot{y} - x)\mathbf{s}_1 + (\ddot{y} + 2\dot{x} - y)\mathbf{s}_2 + \ddot{z}\mathbf{s}_3 \quad (2.8)$$

In this formulation, $\dot{\omega} = 0$, as the CR3BP assumes circular motion [26]. Meanwhile, the gravitational acceleration of the satellite due to the primaries is

$$\ddot{\mathbf{r}} = -\frac{(1-\mu)\mathbf{r}_1}{r_1^3} - \frac{\mu\mathbf{r}_2}{r_2^3} \quad (2.9)$$

where r_1 and r_2 are the distances from the satellite to the larger and smaller primary, respectively. At this point, the equations of motion for the satellite in the 3-body problem are obtained (in the synodic rotating frame) as:

$$\ddot{x} - 2\dot{y} - x = -\frac{(1-\mu)(x-\mu)}{r_1^3} - \frac{\mu(x+1-\mu)}{r_2^3} \quad (2.10)$$

$$\ddot{y} - 2\dot{x} - y = -\frac{(1-\mu)y}{r_1^3} - \frac{\mu y}{r_2^3} \quad (2.11)$$

$$\ddot{z} = -\frac{(1-\mu)z}{r_1^3} - \frac{\mu z}{r_2^3} \quad (2.12)$$

as laid out by Wiesel [44].

2.2 Lagrange Points and Stability

The equations of motion (EOMs) of the CR3BP lead to the existence of Lagrange points; where gravitational and rotational forces on the satellite balance each other out. The location of these points is determined by setting $\ddot{x}$, $\ddot{y}$, $\ddot{z}$, $\dot{x}$, $\dot{y}$, and $\dot{z}$ to zero.

This means that all Lagrange points in the rotating frame obey the equations

$$-x = -\frac{(1-\mu)(x-\mu)}{r_1^3} - \frac{\mu(x+1-\mu)}{r_2^3} \quad (2.13)$$

$$-y = -\frac{(1-\mu)y}{r_1^3} - \frac{\mu y}{r_2^3} \quad (2.14)$$

$$0 = -\frac{(1-\mu)z}{r_1^3} - \frac{\mu z}{r_2^3} \quad (2.15)$$

as presented in Wiesel [44]. From these, one can rapidly gather multiple pieces of information about the Lagrange points; the first being that $z = 0$, per Equation 2.15, meaning that all Lagrange points lie in the orbital plane of the primaries. One can also gather that $r_1 = r_2 = 1$ will represent a pair of Lagrange points (L4 and L5), in which the top two equations simplify to

$$-x = -x + \mu + \mu x - \mu^2 - \mu x - \mu + \mu^2 = -x \quad (2.16)$$

and

$$-y = -y = \mu y - \mu y = -y \quad (2.17)$$

and z is forced to zero. Because $r_1 = r_2 = 1$ where 1 is the distance between the centers of mass of the primaries, these points lie at the vertices of two equilateral triangles; each one having the primaries as the remaining vertices. The remaining equilibrium points lie on the line thru the primaries obtained by setting both y and z to zero. This results in a quintic equation for x , which typically has no more than three real roots [44]. For a system with masses that are roughly the same relative order of magnitude as the Earth-Moon system, these roots lie between the primaries (L1), on the far side of the larger primary from the smaller one (L3), and on the far side of the smaller primary from the larger one (L2) [44].




Figure 2. Earth-Moon Lagrange Points: the five Lagrange points in the Earth-Moon system. [Graphic from Wikipedia, modified in Paint]

These Lagrange points are of interest because satellites can orbit about them; which opens up a slew of potential new parking orbits for this research. Specifically, Lyapunov orbits about the L1 and L2 points have been included as nodes. The application of these orbits to military purposes has been studied extensively by Brick [3] and Ostman [35], and many of their findings have been incorporated into this investigation.

2.3 Orbits of Interest

The CR3BP offers a limitless supply of potential novel orbits that could be used for satellites and cargo, but the implementation of a logistics network requires a finite set of transfer paths and parking orbits. The scope of this research focuses on periodic orbits in the rotating frame. As previously mentioned, the nature of the CR3BP means that any orbital predictions will be numerical and approximate [12], [38]. To put this in more practical terms, the prediction of a satellite’s position and velocity will be less accurate farther forward in time [43]. The exception to this is periodic orbits; a set of orbits in cislunar space where, once a full orbital period has been solved for, the position of a satellite in that orbit can be estimated with relative accuracy as a function of time [12], [38]. These orbits have been used in this investigation as the finite set of storage nodes and transfer pathways.


A unique advantage offered by cislunar periodic orbits is that many of them are inherently unstable; objects placed there will drift away over time. This instability, while it does mandate some station-keeping (usually on the order of centimeters per second of delta-V per orbit [3]), allows spacecraft to move from one to another with little propellant expenditure. In logistics terms, this means that a cislunar orbit can be ideal for short-to-mid term storage of service vehicles, replacement satellites, etc., where the low transfer costs would outweigh the higher station-keeping costs [35], [3].

Low energy-transfers are discussed in depth in Ostman [35], but a brief summery is offered here. Mathematically, although energy is not constant in the CR3BP, there exists an energy-like value called Jacobi’s Constant. Zero-Velocity Curves (ZVCs) are used to describe where a satellite in a CR3BP can reach with a (mathematically) real change in velocity vector for a given Jacobi Constant. An example of ZVC diagrams is shown below, from Ostman [35].


The paths between these reachable regions provide an useful logistical feature by requiring little delta-V to enter and exit [35]. The other low-energy maneuvers that cislunar trajectories are particularly useful for are inclination changes with respect to Earth. The delta-V required for an inclination change is given by the formula:

$$\Delta V = 2V_{initial} \sin \frac{\theta}{2} \quad (2.18)$$


[1] where θ is the plane change angle and $V_{initial}$ is the velocity at the start of the plane change. Generally for a given inclination change, the slower the satellite is moving, the cheaper an inclination change will be. This is why it is possible to raise a satellite’s apoapsis (the highest point of the orbit, where velocity is lowest), perform the change at the new apoapsis, and then lower the orbit to achieve a cheaper inclination change [1]. Since cislunar orbits are inherently quite far from Earth compared to LEO thru GEO orbits, they offer opportunities for relatively low-cost inclination changes with respect to Earth, at least compared with the cost of performing such maneuvers in the LEO-GEO volume. The feasibility of taking advantage of this phenomena is the reason for the inclusion of a hypothetical polar GEO constellation in this investigation.




(a) $JC = 3.2551$



(b) $JC = 3.1862$



(c) $JC = 3.0926$



(d) $JC = 3.0035$

Figure 3. Zero Velocity Curves: an example of ZVCs for four different Jacobi constants. The grey regions are called the “forbidden regions”, and represent the areas that cannot be reached by a ZVC. Going below a certain maximum JC is necessary to enable low-cost transit to and from Lagrange points (represented by the plus symbols). [Image credit: Ostman [35]]

2.4 Generalized Multi-Commodity Network Flows

Modeling of logistics networks in the interplanetary space environment first came about in 2006, largely in response to a national prioritization of space exploration [40]. Hence, these networks are a fairly recent field of study with limited variation in methods. The method used in this investigation is a Generalized Multi-Commodity Network Flow (GMCNF) formulation, as originally developed by Ishimatsu [18]. This method has subsequently been refined into Time-Expanded and Event-Driven versions, largely thanks to the efforts of Ho [21]. These TE-GMCNF and ED-GMCNF models address the key issue with the first-proposed static networks: they can account for the passage of time. The difference between TE and ED networks is subtle, but mathematically important. This difference is best illustrated in Figure 4. In such flows, commodities move along arcs between nodes. The commodities include not just items such as payloads and propellants, but everything that physically moves through the network such as the structure of the satellites and transfer vehicles. Nodes represent physical locations of interest, such as launch sites or periodic orbits. Arcs represent the pathways that commodities can take to travel between nodes. It is important to note that there can be multiple arcs between the same pair of nodes (i.e. if one orbit transfer was optimized for speed, and another was optimized for propellant). There can also be holdover arcs that loop back to their node of origin. These represent processes such as a satellite going around its parking orbit and returning to the insertion node, or a factory consuming resources over time to build satellites [19].

A Time-Expanded GMCNF (TE-GMCNF) as defined in Ho [22] is simply the addition of a time dimension to the static GMCNF. If the static GMCNF is a 2D diagram of nodes and arcs, then the TE-GMCNF is a 3D structure composed of layers of GMCNFs. Each arc in the TE-GMCNF now runs from one node in one discrete moment in time, to another node in another moment in time. A holdover




Figure 4. TE-GMCNF and ED-GMCNF Examples: note how the movement arcs on the TE version encompass both physical movement and the passage of time; while these are separate on the ED version. [Generated in PowerPoint]

arc can be envisioned as a vertical line connecting the same node in multiple times. How the discrete moments in time are selected remains an area of active research; Ho proposes an event-based time step model [22] to account for low-thrust trajectories, the time-of-flight (TOF) of which is not known beforehand.

An Event-Driven GMCNF (ED-GMCNF) exchanges the fixed time steps between event layers for a variable-length time step. It does this by restricting movement through space to movement arcs and movement through time to hold arcs, as shown in Figure 4. This allows the optimizer to calculate the duration of each event layer while the network is being optimized. Allowing the length of time between layers to be calculated on the fly, as opposed to having static time steps, is critical because low-thrust trajectories must have their transfer time calculated at the moment of departure, based on current vehicle mass. It is important to note that the relationships between mass and TOF for vehicles on low thrust arcs is non-linear; however the

problem can be kept linear by fitting a piecewise linear function to the curve [21]. The accuracy of this method is limited only by the computation power available. Although low thrust trajectories ultimately were not used in this investigation, the ED-GMCNF was still selected for testing because of this capability; the improvements and findings of this research do not preclude the use of low-thrust, and adding it is a logical next step.

The ED-GMCNF in this investigation includes the following elements:

- A set of all transportation arcs $\mathcal{A}_T$
- A set of all hold arcs $\mathcal{A}_H$
- A set of all nodes $\mathcal{N}$
- A set of all event layers $\mathcal{E}$
- A set of all vehicles $\mathcal{V}$

The independent variables in the ED-GMCNF of this investigation are the commodity vectors x_{ijvt}^+ and x_{iie}^- . In these variables:

- i and j and the “from” and “to” nodes, respectively
- v is the vehicle
- t is the type of arc (fuel-optimized, time-optimized, or launch)
- e is the event layer

In terms of an optimization/linear programming problem, a ED-GMCNF attempts to minimize the cost function

$$\mathcal{J} = \sum_{e \in \mathcal{E}} \sum_{(i,j,v,t,e) \in \mathcal{A}_T} c_{ijvt}^+ x_{ijvt}^+ + \sum_{e \in \mathcal{E}} \sum_{(i,i,e) \in \mathcal{A}_H} c_{ii}^+ x_{ii}^+ \quad (2.19)$$

Over the course of this investigation, it was determined that the most realistic set of cost coefficients was 1 for all propellant commodities on movement arcs, and 0 for everything else. This produced a network that would minimize fuel usage while remaining within time limits.

The ED-GMCNF is subject to the demand constraint

$$\sum_{(j,v):(i,j,v,t) \in \mathcal{A}_T} x_{ijvte}^+ - \sum_{(j,v):(j,i,v,t) \in \mathcal{A}_T} x_{jivte}^- + x_{iie}^+ - x_{iie}^- \leq d_{ie} \quad \forall i \in \mathcal{N} \quad \forall e \in \mathcal{E} \quad (2.20)$$

where d_{ie} is the demand (negative) or supply (positive) at node i at event e . The sum of all commodities arriving minus the sum of all commodities leaving must be less than or equal to the demand or supply.

The network is also subject to the transformation constraints

$$\mathbf{B}_{ijvt} x_{ijvte}^+ = x_{ijvte}^- \quad \forall (i, j, v, t, e) \in \mathcal{A}_T \quad \forall e \in \mathcal{E} \quad (2.21)$$

$$\mathbf{B}_{ii} x_{iie}^+ = x_{iie}^- \quad \forall (i, i, e) \in \mathcal{A}_H \quad \forall e \in \mathcal{E} \quad (2.22)$$

which enforce the consumption of fuel by relating the commodities starting an arc (x^+) with those finishing an arc (x^-). The B matrix is derived from the ideal rocket equation.

The concurrency constraints

$$\mathbf{C}_{ijv}^+ x_{ijvte}^+ \leq s_{ijv}^+ \quad \forall (i, j, v, t, e) \in \mathcal{A}_T \quad \forall e \in \mathcal{E} \quad (2.23)$$

$$\mathbf{C}_{ii}^+ x_{iie}^+ \leq s_{ii}^+ \quad \forall (i, i, e) \in \mathcal{A}_H \quad \forall e \in \mathcal{E} \quad (2.24)$$

enforce the fact that some commodities must travel with other commodities, such as fuel requiring a tank. These matrices also enforce miscellaneous constraints, such as the fact that the service vehicle cannot travel on launch arcs to the Earth and Moon's surfaces.

The non-negativity constraints

$$x_{ijvte}^{\pm} \geq 0_{k*1} \quad \forall (i, j, v, t, e) \in \mathcal{A}_T \quad \forall e \in \mathcal{E} \quad (2.25)$$

$$x_{iie}^{\pm} \geq 0_{k*1} \quad \forall (i, i, e) \in \mathcal{A}_H \quad \forall e \in \mathcal{E} \quad (2.26)$$

prevent any commodities from having values below zero. This means that all vehicles must begin and end their movement arcs with fuel in their tanks, as well as generally preventing situations that could not exist in reality.

Finally, the mass and time constraints

$$y_{ijvte}^{\pm} = \mathbf{M}x_{ijvte}^{\pm} \quad \forall (i, j, v, t, e) \in \mathcal{A}_T \quad \forall e \in \mathcal{E} \quad (2.27)$$

$$f(x_{ijvte}^+, y_{ijvte}^+) \leq t \quad \forall v \in \mathcal{V}' \subseteq \mathcal{V} \quad \forall e \in \mathcal{E}' \subseteq \mathcal{E} \quad (2.28)$$

allow for manipulation of data via the masses (rather than quantities) in the commodity vector, and to ensure that event layers are completed with their time limits.

Compared to previous work by Ho [21], the number of independent variables has been reduced to solely include x_{ijvte}^+ and x_{iie}^+ , as all other variables are dependent on them. For example, whenever x_{ijvte}^- is required, $B_{ijv}x_{ijve}^+$ can be substituted.

2.5 Mixed Integer Linear Programming

The principal tool for solving a network optimization problem is mixed integer linear programming. Any network in which n resources factor into the cost function can be modeled as a constrained $n + 1$ -dimensional polytope (similar to a hyperplane with flat sides) [28]. The optimal solution will then correspond to a minimum or maximum point on this polytope. For example, if the objective function had two variables, the polytope would be a constrained 3D surface with the x and y axes as the input variables, and the z axis as the cost to be minimized. An example is shown in Figure 5.

The “mixed integer” part of the problem is a feature common to logistics networks due to the fact that certain resources must be constrained to integer values, such as the number of satellites transported in a single launch. This necessitates a more complicated solution method than for a normal (non-integer) optimization problem. By default, a mixed integer linear problem (MILP) will have a convex solution space, which significantly simplifies the solution process [32]. A common, but crude, solution, is to solve the MILP problem as an LP problem, and then round off the numbers that are supposed to be integers. However this method tends to produce results outside the solution space; even more so in high-dimensional solution spaces [32]. An alternative is called “branch and bound”, which begins in the same manner as the crude solution above; solve the comparatively-easy LP problem. If all the commodity values that need to be integers are, then the solution has already been found. However, if they are not, rather than simply rounding, the branch-and-bound method reruns the problem with a new constraint, namely that the misbehaving quantity q , which should be an integer but is instead a non-integer r , must now obey the constraint

$$q \leq \lfloor r \rfloor \tag{2.29}$$




Figure 5. Linear Programming Cost Function: graphical representation of a linear programming solution with two variables (y and z) feeding into a cost function (x). The optimal solution is the lowest point. [Image credit: Wordpress]

meaning that q must be less than or equal to the first integer below r , or the constraint

$$q \geq \lceil r \rceil \tag{2.30}$$

meaning that q must be greater than or equal to the first integer above r .

In this investigation, the MILP optimization of the ED-GMCNF was handled by Gurobi’s MATLAB plugin. Gurobi utilized the above branch and bound method to arrive at an optimal solution [16]. In the ED-GMCNF, all constraints are linear: the rocket equation is solved for all impulsive arcs before the transformation matrices are handed to the optimizer, and a similar tactic is used for low-thrust arcs (when included) by breaking their mass/delta-V curves into piecewise linear functions [21].

2.6 Sparse Non-Linear Trajectory Optimization

Although not the focus of this investigation, a logistics network can only be optimized properly if it is provided with consistent data. The optimization of orbits for various purposes has existed in one form or another since the advent of space-flight [36], and this investigation makes use of the latest advances in optimization for the purposes of minimizing both fuel usage and TOF. Optimization techniques can broadly be divided into direct and indirect methods; however indirect methods have largely fallen out of favor in recent years due to their lack of robustness; requiring more accurate initial guesses than direct methods [37]. The primary tool used in this research was the Sparse Non-Linear OPTimizer provided by STK’s Astrogator.

The Sparse Non-Linear OPTimizer is an optimizer that uses sequential quadratic programming (SQP) to find locally optimal solutions [15]. SNOPT is a robust tool that can cover a wide range of solutions, but still requires a decent initial guess to be effective (i.e. to find the *correct* local optimum. Combining SNOPT with STK’s

differential corrector and basic trajectory design best practices (burn at apoapsis or periapsis, perform inclination changes at max. radius, etc.) enabled the quick generation of numerous consistently-well-designed trajectories to test the ED-GMCNF formulation. As shown in Chapter III, these trajectories were comparable in quality to the much more exhaustively designed cislunar trajectories of Brick and Ostman [3] [35].

2.7 Literature Review

As mentioned in Chapter I, the concept of space logistics networks is a relatively recent concept, though the broader concept of a mathematically optimizable logistics network is not. The space-domain concept was born as a result of a 2004 national security-focused push for space exploration, and initial models consisted of the same time-expanded networks that had previously been used for terrestrial applications [40]. Examples of these models include optimization of road networks for traffic flow [29] and optimization of country-wide water systems for minimum cost and energy use [18]. Development of multi-commodity networks for terrestrial transport has continued since then, and several iterations offer unique features that could be applied to future space-domain networks. Such features include safety rating for arcs [7], maximum arc capacities [7], and the first instances of time-expanded networks [25]. There has also been significant research into the computational time required to solve multi-commodity networks of various sizes [41], but the numbers of nodes and arcs required to make that time a significant obstacle has been shown to be far higher (on the order of thousands) than what is realistic for a space-based network [41], which will by necessity be restricted to periodic or near-periodic orbits [38].

Some work has also been done to effectively solve a multi-commodity network in reverse: given a set of flow requirements, determine the optimal locations for nodes [34].

The aforementioned terrestrial models required frequent and occasionally crude modification, such as imaginary nodes, to reflect the reality of interplanetary logistics [40]. The generalized multi-commodity network flow used in this investigation was developed in 2013 by Ishimatsu [18]. In particular, its inclusion of the flow transformation matrix B to model fuel consumption eliminated the need for imaginary nodes in the middle of orbit transfers. The multi-commodity flow problem is not itself recent, appearing as early as the 1970's [45] [29], but its application to interplanetary logistics is. This is evidenced by its lack of inclusion in a 2018 survey of current multi-commodity network flow applications [41].

Research into the circular restricted 3-body problem has focused around finding orbits of interest to serve as nodes for the GMCNF. A logistics network of this type does not include trajectory optimization by nature; the values of transfer time, delta-V, etc. are fixed for each arc and an optimal solution simply picks the best ones. Consequently, the exploration of Brick into the military applications of the CR3BP[3] was invaluable for its detail descriptions of low-cost orbit transfers, as was the work of Dahlke [12].

Much like the CR3BP, linear programming, specifically mixed-linear programming, has been a well-documented field of study since the 1960's [28]. Also like the CR3BP, research for this investigation was done in support of generating the GMCNF models, rather than in an attempt to innovate new methods. Itai [20] showed in 1978 that a two-node network was analogous to a linear programming problem. More recently, Ishimatsu offers an explanation of how GMCNFs can be solved with linear programming [18] that is better geared to the space-domain, and LeValley's work

provided excellent insight into the best tools available at the time of this writing for solving these problems [31]. A piece of software called *SpaceNet* was investigated as well. This software was developed by MIT for the express purpose of solving space logistics networks, and was instrumental in the investigations of Ho [22], [17]. However, this software has not been updated since 2012, and does not use the GMCNF.

2.8 Contributions of this Research

This research has made two primary contributions to the work by Ho [21] that it is based on: a significant reduction in independent variables and the inclusion of arc types as an additional option. The former was achieved by making the x_{ijvte}^- and x_{iie}^- variables the only independent variables. As shown in Equations 2.22 and 2.27, all other variables can be modeled on these. This feature significantly reduced the mathematical complexity of the network compared to Ho’s version. The arc type designator t was added to enable the network optimizer to select from fuel and time-optimized arcs between the same node pair, and did not exist in previous iterations.

In addition to the above, this research, specifically the test network, represents the first use of this ED-GMCNF formulation to model a reusable launch vehicle.

2.9 Summary

This chapter has covered the mathematical basis for this research, including the CR3BP, Mutlicommodity networks, and MILP problems. This chapter also included a summary of past work in the field, and outlined the contributions of this research to the existing body of work.

III. Methodology

This chapter of the investigation contains an overview of how the investigation built on the background data to address the problem statement. This chapter begins with an overview of the assumptions and equations that were used in the GMCNF formulation, and a discussion of the proof-of-concept test network. Following that, this section covers the different trial networks that were explored and the reasoning behind them.

The methodology herein was developed to answer the research questions described in section 1.3. The nodes and arcs used were drawn from the work of Brick and Ostman [3], [35]. The delta-V and TOF data was either derived from the aforementioned sources or calculated using STK’s Sparse Nonlinear OPTimizer (SNOPT) and differential corrector tools. A test problem was developed and simulated as a proof of concept.

3.1 Assumptions and Equations

This section includes the equations and assumptions common to all of the networks explored in this investigation, and the reasoning behind them.

3.1.1 Propellant Mass Fraction.

Propellant consumption and production rates are key to establishing the flow transformation B matrices in the GMCNF model.

The propellant mass fraction is taken from the rocket equation and describes the amount of propellant needed to produce a change in velocity. This is identical to the assumption made in Ishimatsu [18]. Since a delta-V is an inherent component of arcs in a space-based logistics network (albeit delta-V might have a value of zero for a

hold arc), a propellant mass fraction can be defined for any combination of propulsion system and arc from the equation

$$\phi_{ij} = 1 - \exp\left(-\frac{\Delta V_{ij}}{I_{sp}g_0}\right) \quad (3.1)$$

Where ΔV_{ij} is the change in velocity inherent to arc (i, j) , I_{sp} is the specific impulse inherent to the propulsion system, and g_0 is the natural gravitational acceleration ($9.81m/s^2$). The propellant mass fraction ϕ_{ij} is also given by

$$\phi_{ij} = \frac{\Delta m}{m_s + m_{pr}^+} \quad (3.2)$$

where Δm is the change in mass due to consuming propellant, m_s is the structural/dry mass of the spacecraft, and m_{pr}^+ is the propellant mass at the start of the burn,

$$m_{pr}^- = m_{pr}^+(1 - \phi_{ij}) - \phi_{ij}m_s \quad (3.3)$$

3.1.2 Inert Mass Fraction.

Related to the propellant mass fraction is the inert mass fraction, which is part of the **C** matrix; establishing a requirement for the amount of propellant-storing commodities needed to haul a certain amount of propellant. This is an inherent characteristic of the transfer vehicle's structure, and can be reduced by reducing vehicle mass. The equations presented here are taken from Ishimatsu [18]

$$f_{\text{inert}} = \frac{m_s}{m_s + m_{pr}^+} \quad (3.4)$$

where f is the inert mass fraction, m_s is the structural mass, and m_{pr}^+ is the starting propellant mass. This establishes a ratio of structure to propellant mass of

$$\eta = \frac{m_s}{m_{pr}^+} = \frac{f_{\text{inert}}}{1 - f_{\text{inert}}} \quad (3.5)$$

3.2 Transformation Matrix

In an identical manner to Ho [21], the ideal rocket equation was used to find the mass transformation constant 1p for each arc

$${}^1p = e^{\frac{-\Delta V}{g_0 * Isp}} \quad (3.6)$$

where g_0 is the universal gravity constant, Isp is the specific impulse of the vehicle, and ΔV is the delta-V for the arc. The inert and propellant mass fractions were combined with this quantity 1p to form the transformation **B** matrices for both the test and main networks. A sample one is

$$\mathbf{B} = \mathbf{M}^{-1}(\mathbf{QM} + \mathbf{R}) \quad (3.7)$$

where $\mathbf{R} = 0$ when the rocket equation is satisfied (this matrix would be non-zero if there were a set fuel cost associated with travelling on an arc, rather than a purely delta-V dependent one). Also in equation 3.7, $\mathbf{M}$ is the mass transformation matrix, and

$$\mathbf{Q} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ {}^1p - 1 & {}^1p - 1 & {}^1p & {}^1p - 1 & {}^1p - 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.8)$$

Note that in this equation, the propellant is the third commodity out of five.

3.3 Proof of Concept

This investigation began with the optimization of a “test network”, a simple network with 4 nodes, 7 arcs, and 3 commodities shown in Figure 6. This test network would provide a foundation for subsequent, more complex, problems. This test network also assisted in the selection of programming languages and methods. The test network considered contained the nodes shown in Figure 6 and listed in Table 1.

This network was modelled using the mixed-integer linear programming techniques outlined in the work of Ho [22]. This involves modeling commodities in a state vector, which is multiplied by matrices to reflect changes due to the passage of time and/or the expenditure of resources [22]. This method is detailed in section 2.4. To simplify this test network, only a single type of vehicle was considered: a hypothetical two-stage rocket. The first stage would be capable of carrying the second stage and payload to LEO, and would be able to land back at KSC for reuse. The second stage would be capable of carrying the payload and itself to a safe landing at the ISRU station, at which point the second stage would be refueled and function as a tug for completed satellites. At the end of the second stage’s lifespan, the tug would be put into a graveyard orbit or de-orbited depending on where it reached the end of its life. In addition, the test network makes the assumption that the constellation being

Table 1. Test Network Nodes: the nodes included in the test network, and their purposes.

Node	Purpose
Earth launch site	Starting point for all commodities
Stage separation point	Node to allow lower stage to separate
L2 Lagrange point parking orbit	Satellite storage
GEO constellation	Hypothetical constellation under service

serviced lies at geosynchronous altitude in the moon's orbit plane; this means that the delta-V required to transfer between the L2 parking orbit and the constellation remains relatively constant over time. This test network involves the following arcs:

- Transportation arc from KSC to LEO, representing the launch of replacement satellites.
- Transportation arc from LEO to KSC, for return of the first stage.
- Transportation arc from LEO to L2, for insertion of replacement satellites into the parking orbit.
- Transportation arc from L2 to GEO, representing the replacement of a satellite.
- Transportation arc from GEO to L2, representing the return of the upper stage/service vehicle.
- Transportation arc directly from LEO to GEO, offering an alternative route to the constellation.
- Hold arcs for each node to connect event layers.

The test network is shown in figure 6.

This problem featured the following additional constraints, implemented through the concurrency and arc activation matrices:

- The lower stage and lower stage propellant may only travel on arcs 1 and 2 (to and from the launch pad).
- Cargo may not sit at the launch pad (no holdover arc at node 1)
- Except for the lower stage, commodities may not rest at the stage separation point, as it is not a stable orbit (no holdover arc at node 2)




Figure 6. Test Network: features the reconstitution of a GEO constellation using satellites stored in a halo orbit around the L2 Lagrange point. [Generated in PowerPoint]

The transformation (B) matrices in this problem follow the ideal rocket equation

$$\Delta V = Isp * g_0 * \ln\left(\frac{m_f}{m_e}\right) \quad (3.9)$$

where Isp is the specific impulse of the vehicle, g_0 is Earth’s gravitational acceleration, m_f is the mass of the entire vehicle at the start of its arc ($M \times x_{ijve}^+$), and m_e is the mass of the entire vehicle at the end of the arc ($M \times x_{ijve}^-$). Note that satellites, unused propellant, and upper stage/service vehicle (when applicable) are all lumped into the vehicle’s empty mass.

3.4 Test Problem Results

The test network described in Section 3.3 was run in five configurations. The first was a test run without the “shortcut” arc. The results of this are shown in figures 7 thru 10, each representing one event layer.

Of the other four runs of the test network, the first two were allowed 2 (simulated) days to complete, and consisted of a trial with weight in the objective function assigned only to the structures moving (i.e. the optimizer would attempt to complete the network with as few movements of upper and lower stages and satellites as possible), and a trial with the weight assigned to fuel (i.e. the optimizer will attempt to move as few kg of fuel as possible). The second two trials mimicked these, but were allowed 10 (simulated) days time to complete. All computations were run on a Lenovo Y50-70 laptop with an Intel i7-4720HQ 2.60 GHz processor and 16 GB of RAM, and none took more than 0.1 seconds to complete (the average for the test network was 0.08 seconds). The results of these four trials are presented below. As a disclaimer, the numbers used for delta-V and TOF were designed to test the effects of changing parameters on the optimizer, and were not based on actual trajectories.




Figure 9. Test Problem 3rd Event Layer: the demand for four spare satellites at the parking orbit has been met, so they disappear. [Generated in Microsoft PowerPoint]




Figure 10. Test Problem 4th Event Layer: the service vehicle makes the round trip in one step. The demand for a replacement satellite is satisfied, so it disappears. [Generated in Microsoft PowerPoint]

Subsequent investigations utilize figures derived from STK-plotted trajectories and actual launch systems and spacecraft.

The obvious choice for the optimizer is to utilize arc 6, the "shortcut" arc from the LEO stage separation point to the GEO constellation, as this arc has been assigned a delta-V of only 1 km/s, as opposed to a total of 4.9 km/s to pass through the parking orbit on the way there (arcs 3 and 4). It has also been assigned a TOF of just one day, as opposed to 4 days to pass through arcs 3 and 4.

The optimizer takes the shortcut arc whenever instructed to conserve fuel, regardless of TOF. When the optimizer is instructed to minimize the number of structures moved (satellites and lower and upper stages), however, it attempts to take the "long way round" when not forced to use the shortcut by a limited time. This is because there is a demand for satellites in the parking orbit as well as in the constellation, and by taking this route it can drop them off on the way there, thus reducing the number of objects that have to be moved. This is consistent with the expected behavior of the optimizer.

There are a number of other considerations that came to light during the test network trials, mostly regarding holdover arcs and the passage of time. In an ED-GMCNF, time only passes for a commodity when it uses a holdover arc. As mentioned in Section 2.4, this disconnection between physical and temporal movement is what makes the use of low-thrust trajectories possible. However, it also makes it a bit harder to synchronize the results of the optimizer with real-life movement of assets. The biggest example in the test network is the reusable lower stage. Because it takes time for the lower stage to carry its cargo to the stage separation point and return, it and its propellant must pass through a holdover arc at some point in their journey. However, the only two nodes they will visit are the stage separation point and the launch pad. The former is not an orbit, but a point in the upper atmosphere,




Figure 11. Test Network with “Shortcut” Arc: the test network network including a “shortcut arc” directly from the stage separation point to the constellation being serviced.

anything attempting to “wait” there in real life would crash back down to the Earth. The latter is fine for waiting physically, but if commodities are allowed to use this holdover arc, it will create problems in more complex networks with multiple launch windows, for example as the optimizer erroneously adds commodities from a previous launch to a new one. Ultimately, both of these problems can be corrected by only allowing certain commodities (i.e. the lower stage and its propellant) to use these arcs at certain event layers, but it must still be taken into careful consideration.

The reason that the lower stage in the test network has been forced to hold at the stage separation point rather than at the launch pad is because, as stated in Section 2.4, the length of a event layer in an ED-GMCNF is the same as the length of time as the longest sum of vehicle movement arcs within that event layer. This is not an issue in the test network, as, after the stage separation point, the upper stage must make a much longer journey out to GEO or L2 than the lower stage must make back

down to the pad. However, if the upper stage had a very short arc to travel, such as a circularizing burn, allowing the lower stage to return in the same event layer would make the results appear as if the circularizing burn took much longer than it did, since the launch and return arcs would be longer. In virtually any real-world scenario, mission planners will be far more concerned with the satellites getting into position on time, as opposed to having the lower stage return on time. Hence, the decision has been made to make the lower stage wait at the stage separation point. All other commodities are forbidden from using that node as a holdover in order to reflect its instability as best as possible.

One other limitation of the ED-GMCNF demonstrated here is that TOF is not a quantity that the ED-GMCNF can directly optimize for, since it is a constraint and not part of the objective function. It is possible to simply iterate the network with lower and lower “maximum TOF allowed” values until it returns as unsolvable, but this is a fairly crude method. In future trials, when calculated delta-V and TOF values are provided (rather than the estimates here), fuel-optimized arcs are calculated beforehand, and commodities are forced to use them in the network. As seen above, this is necessary because attempting to minimize fuel use will almost certainly preclude the use of minimum TOF arcs (a slower, more efficient arc would be used if available), and attempting to optimize for non-fuel commodities will produce nonsensical results.

3.5 Main Network

The main body of this investigation consisted of a single network, similar but not identical to the toy network, which was run in various configurations. For the purpose of expediency, this was accomplished by selectively deactivating unwanted nodes and

arcs (prohibiting commodities from using them), rather than by removing them from the network entirely. The effect on results was identical.

3.5.1 Nodes and Arcs.

The main network for this investigation was designed to leverage previous work on cislunar trajectory optimization to provide a complete diagram of potential orbits and transfers, with multiple routes for commodities to “choose” to reach the constellations. The primary driving factor behind the exact nodes and arcs chosen for use was the cislunar trajectory optimization work of Brick [3] and Ostman [35]. Leveraging data from these works on periodic orbits and impulsive transfers in the cislunar domain saved time and helped minimize the likelihood of the network optimizer returning misleading results, since a poorly optimized trajectory could be rejected for a mission when a better-optimized one would have been the best solution. The work of Dahlke [12] focused on low-thrust trajectories, which were not included in this investigation.

3.5.2 Delta-V and Time of Flight Information.

The numbers used for delta-V and TOF in the main network were taken, as much as was possible, from the works of Brick [3] and Ostman [35], which have been extensively optimized to take advantage of low-energy pathways between orbits using impulsive maneuvers. Unfortunately, the combined data of these two sources is not sufficient to create a complete Earth-to-constellation service network. The gaps in data (transfers such as LEO-to-L1) not covered by these works were modelled using AGI’s Astrogator module for STK. This modelling used the default cislunar propagator available in that module and the SNOPT and differential corrector tools to create trajectories optimized for both delta-V and for TOF. Although not as rigorous as the optimization techniques presented in Brick, Ostman, and Dahlke, consistency

was maintained by using their orbits as starting points for the transfers. The SNOPT optimizer is the same one that was used by this previous work, and when tested directly against it, it produces results that lie on the generated Pareto fronts. For example, the STK SNOPT, when set to minimize delta-V on a transfer from L2 Lyapunov to GEO, produced a delta-V of 18 km/s including both impulsive burns, and a TOF of 6.64 days. Ostman’s Pareto front for the same transfer places a 6.70 day burn at 16.7 km/s delta-V. This is shown graphically in Figure 14. While not quite Pareto efficient; these results are similar enough as to not affect which arcs are chosen. Some discretion is necessary here; if delta-V were unlimited, a straight burn to a new orbit would be the minimal time solution; but this is obviously not practical. Again, this can be seen in Figure 14, with massively diminishing returns towards either end of the Pareto front. This has serious implications for the investigation; if an arc were modelled with two impractical trajectories (i.e. a time-optimized trajectory requiring too much delta-V and a fuel optimized trajectory taking too long to arrive), it would never be selected even if a trajectory closer to the Pareto front’s knee point [30] would be the best option. Consequently, every effort has been made to keep the trajectories used a consistent distance from the knee point. For trajectories generated using STK, this was accomplished in one of two ways. First, by imposing bounds on burn angle and total delta-V and using SNOPT. An example of this is the L2 Lyapunov to LLO arc shown in Figure 12. Alternatively, a delta-V consistent with literature values and similar trajectories was selected for the injection maneuver, and STK’s differential corrector was used to calculate the proper burn time and angle. This technique was used to replicate NASA’s estimates for high and low delta-V transfers between the Earth and Moon [10]. An example of a trajectory generated this way is the LLO to GEO arc shown in Figure 13.




Figure 12. L2 Lyapunov to LLO Transfer: this transfer was generated using SNOPT for the injection burn and differential corrector for the arrival burn. [Graphic generated using STK]




Figure 13. LLO to GEO Transfer: this transfer was generated using the differential corrector for both the injection burn timing and the arrival burn. [Graphic generated using STK]

For trajectories leaving the lunar plane, i.e. those to the hypothetical polar constellation, the inclination change was modeled as occurring entirely at the highest burn point. No intermediate points or partial inclination changes were used. In both time and fuel-optimized trajectories, this change was computed with an objective of minimizing delta-V, since all burns are assumed to be impulsive.

In order to simplify the investigation, all trajectories were assumed to be reversible (i.e. travelling from node 1 to node 2 will cost the same as travelling from node 2 to node 1). This assumption discounts potential aerobraking and lunar gravity assists that could aid transfer in one direction. This is supported by Whitley [42], in which the inclusion of a lunar gravity assist for transfer to a 9:2 near rectilinear halo orbit (NRHO) differs by only 9 m/s on the inbound vs outbound transfer and insertion maneuvers (0.404 vs. 0.395 km/s). This is only a 2% difference and not nearly enough that an orbit would be erroneously discarded or included by the network optimizer.

Finally, the decision was made to omit any low thrust trajectories from the main network. The ED-GMCNF is specially formulated to allow them [21] and future work will undoubtedly make use of them, so the ED-GMCNF is retained as the type of network under test. The other reason for this omission is that lunar ISRU and refueling is a significant portion of the main trials performed, and the Moon’s geology is not conducive to the synthesis of low thrust propellants [27], while lunar ice is conducive to the production of hydrogen-based propellants.

3.6 Main Network Parameters

Based on the limitations of the test network, the main network was constructed with the nodes in Table 2 and the commodities in Table 3. The values of delta-V and TOF required for transport between these nodes can be viewed in Appendix A.




Figure 14. Example Pareto Front: Pareto fronts from Ostman for transfers from the L1 and L2 Lyapunov orbits to GEO. The green dot represents the result achieved with a minimum delta-V solution in STK's SNOPT. [Modified from [35]]

As mentioned in the introduction and illustrated in Aerospace’s current listing of R3 missions [24], there are unfortunately no currently operating vehicles designed with extended, multi-thousand meter per second delta-V missions in mind. Characteristics such as structural mass, propellant and payload capacity, and engine specific impulse are realistic estimates, but should not be thought of as representing any existing spacecraft. For all tests, the characteristics in Table 3 were used.

The following assumptions were made during the selection of these nodes:

- All orbits except for the polar orbit lie in the Moon’s orbit plane.
- LEO is at a 200 km altitude.
- LLO is at a 100 km altitude.
- Geosynchronous orbits are at a 35,786 km altitude.
- The L1, L2, and DPO orbits are from Ostman [35].
- Launch from Earth to LEO requires 9.40 km/s delta-V and 0.02 days TOF.
This leaves nearly 200,000 kg of fuel for service vehicle use.
- Launch from Moon to LLO requires 1.73 km/s delta-V and 0.02 days TOF.
This leaves nearly 2.8 million kg of fuel for service vehicle use.

It is critical to note that this network is not designed with replacing broken parts in mind; any damage that would require a replacement part is assumed to necessitate replacement of the satellite. The commodity “spare part” is placed into the network to mitigate a weakness discussed in Chapter V, which renders demand for the service vehicle itself impractical.

The cost function for the main network assigns a coefficient of 1 to all propellant travelling on movement arcs. For propellant on holdover arcs, as well as all other

Table 2. Main Network Nodes: the nodes included in the main network, and their purposes.

Node	Purpose
Low earth orbit	Starting point for service vehicle and satellites
Earth ground station	Starting point for Earth-based refueling mission
Lunar ground station	Starting point for Moon-based refueling mission
L1 Lyapunov Orbit	Potential parking orbit
L2 Lyapunov Orbit	Potential parking orbit
Distant Prograde Orbit	Potential parking orbit
Low Lunar Orbit (LLO)	Potential parking orbit
Geosynchronous Orbit	Hypothetical constellation
Polar Geosynchronous Orbit	Hypothetical constellation

Table 3. Main Network Commodities: the commodities used in the main network.

Commodity	Mass (kg)	ISP (s)	Capacity
Satellites	100	N/A	N/A
Service Vehicle	290	500	10 satellites 5000 kg propellant 100 spare parts
Propellant	N/A	N/A	N/A
Launch Vehicle	10,000	300	5,000,000 kg propellant
Spare Part	0.1	N/A	N/A

commodities, the cost is 0. This is based on the results of the test network and is most reflective of real-life mission design.

The constraints imposed on the main network are as follows:

- The service and launch vehicles must have propellant to use a movement arc.
- Propellant must travel in a vehicle or satellite, and cannot exceed their capacities.
- A satellite cannot use a holdover arc in a constellation node unless a service vehicle is present. This is to prevent the service vehicle from dropping off satellites before they are needed, as a mission planner would not have foreknowledge of when reconstitution would be necessary.
- Spare parts cannot use and arcs without a service vehicle present. Since these are a way to simulate a service (repair), rather than a real item, they cannot be allowed to be separate from the service vehicle.

An example of a concurrency matrix is

$$\mathbf{C} = \begin{bmatrix} -pc1 & -pc2 & 1 & -pc3 & 0 \\ 1 & -sc & 0 & 0 & 0 \\ 0 & 1 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (3.10)$$

this matrix is for a five commodity network, with the commodities being satellites (1), service vehicles (2), propellant (3), launch vehicles (4), and spare parts (5). The first line enforces fuel capacities for three different vehicles. The second line enforces the satellite capacity of a service vehicle. The third line means the vehicles that can move (not satellites) must have fuel to do so. The last line is unused. This matrix

forms the inequality constraint

$$\mathbf{C}\mathbf{N} \leq 0 \tag{3.11}$$

where $\mathbf{N}$ is the commodity vector.

3.7 Methodology Summary

The methodology above was developed to answer the research questions. The nodes and arcs used were drawn from the work of Brick and Ostman [3], [35]. The delta-V and TOF data was either derived from the aforementioned sources or calculated using STK's SNOPT and differential corrector tools. A test problem was developed and simulated as a proof of concept.

IV. Results and Analysis

This chapter describes the results of the main network trials and offers an interpretation of the data gathered. The trials were conducted in four phases:

- With service vehicle refueling allowed from Earth only
- With refueling allowed from both the Earth and moon
- With an additional constellation in an L1 Lyapunov orbit
- With variations to the optimizer tolerances and satellite mass

Unless otherwise noted, the Gurobi optimizer parameters were left at their default values, save for the MIP Gap tolerance (the overall solution tolerance) and the Integer Feasibility Tolerance, which were set to 1e-08 and 1e-09, respectively. These settings were shown to produce networks reflective of reality with acceptable optimization times. All trials also made use of the same number of event layers, limited to a maximum of 15 days apiece unless otherwise noted, and the same commodity demands. These demands were designed to reflect an intense and varied R3 mission supporting a pair of constellations, one at GEO altitude in the lunar plane (Constellation 1) and one at GEO altitude with an 80 degree inclination (Constellation 2). A demand for a repair visit is modelled in network as a demand for the low-mass “spare part” commodity.

In all trials, the fuel launchers/depots had the ability to travel to any node, while the service vehicles could only travel to nodes in space (i.e. not the Earth or Moon surface nodes). In addition, the launchers were required to launch with full tanks; they could not simply dump fuel that would not be needed in the network. Unless otherwise noted, all spare satellites and the service vehicle started in LEO, simulating a launch from Earth.

Table 4. Main Network Demands: the commodities required in the main network.

Event Layer	Demand	Location
2	1 satellite	Constellation 1
4	1 satellite 20 kg propellant	Constellation 2
6	1 repair	Constellation 1
8	1 satellite	Constellation 2
10	1 satellite 1 repair	Constellation 1
12	1 satellite 20 kg propellant	Constellation 2
14	1 repair	Constellation 1

In all result figures, the purple lines represent arcs that were used to transport commodities between nodes (the orange dots). Each column of each graph represents a different node, and each row is a new event layer (time increases upwards). A 1 over an arc indicates that it is optimized for minimum delta-V, and a 2 indicates optimization for time. 3 is for launch arcs and 99 is for holdover arcs. The nodes have been slightly offset vertically to make it easier to tell where each arc begins and ends.

4.1 Earth-Based Refueling Only

The first set of trials of the network were conducted with only the Earth-based launcher available to provide fuel to the service vehicle. The first of these made no exceptions to the above parameters: all spare satellites and the service vehicle started in LEO, the launcher started on the Earth’s surface, and each event layer was allowed a maximum of 15 days to complete. The results of this are shown in Figure 15. This network was rerun with the maximum service vehicle fuel capacity doubled to 10,000 kg in order to confirm that it was not running out of fuel. The results were identical to Figure 15. This confirmed that in this instance, it is more cost effective for the service vehicle to carry the minimum fuel needed to accomplish its task and refuel in

between each than it is to carry extra fuel and satellites and linger in the constellation orbits. At the end of this network, the service vehicle transferred thru DPO, which reduced the delta-V of the inclination change from 3.93 km/s to 3.36 km/s. This is not the cheapest path, going thru both DPO and L1 on the way would cut delta-V to 3.26, but that would have violated the 15 day time limit by taking 17.34 days. However, even this maneuver was only performed when the service vehicle “knew” that it would not have to refuel after accomplishing its mission.

The second trial started the spare satellites in DPO instead of LEO, to determine if this cislunar orbit might be viable for storage if the cost of moving the satellites there in the first place was not taken into account. The results of this trial are shown in Figure 16. This trial conclusively showed that in this case, moving the satellites did not help matters. The service vehicle still chose to refuel as often as possible to avoid carrying excess fuel around.

The third trial returned the satellites to LEO at the start, but moved the launcher to LEO. In this manner, this trial eliminated the fuel launch cost from consideration. The results of this are shown in Figure 17. This trial showed that removing the launch cost did not affect the service vehicle’s behavior. This was rerun with the satellites back in DPO at the start, but this changed nothing in that Refueling the service vehicle at every opportunity was still the optimal solution.

The fourth trial replicated the first trial, but with the maximum time for each event layer lowered to 3 days. The results are shown in Figure 18. Other than the service vehicle using a direct inclination change at the end to save time over it’s cheaper route through DPO, this did not change the conclusions from the previous trials. Note the lower computation time as a significant number of arcs with TOFs longer than 3 days can be written off outright by the optimizer.




Figure 15. Earth Fuel Only Trial 1: this network solution was produced when the optimizer could only launch a refueling mission from Earth. It took 7.5 seconds to optimize.




Figure 16. Earth Fuel Only Trial 2: this network solution was produced when the optimizer could only launch a refueling mission from Earth. Spare satellites stard in DPO. It took 6.0 seconds to optimize.




Figure 17. Earth Fuel Only Trial 3: this network solution was produced when the optimizer could only refuel using fuel starting in LEO. It took 4.9 seconds to optimize.




Figure 18. Earth Fuel Only Trial 4: this network solution was produced when the optimizer could only launch a refueling mission from Earth. The maximum event layer was shortened to 3 days. It took 1.8 seconds to optimize.

Overall, the trials conducted with fuel launched only from Earth make a poor case for the practicality of cislunar orbits as parking spaces for an R3 campaign. Although DPO is used in one instance to reduce the delta-V of a transfer, the need for the service vehicle to refuel and the cost of moving the fuel source appear to override that. In addition; when the schedule was compressed, this too went away in favor of a faster, slightly more expensive inclination change. Based on the network inputs for commodities, masses, delta-V, and TOF, the need to service constellations at multiple inclinations does not make the use of cislunar orbits practical for R3.

4.2 Lunar ISRU Included

The next set of trials added a simulation of lunar ISRU, in that a second refueling mission started on the surface of the moon and could launch to LLO.

The first of these trials was identical to the first of the previous set of trials, all parameters were left as stated in the beginning of this chapter. The results are shown in Figure 19. This result demonstrates a clear preference for lunar refueling over Earth-launched refueling in this scenario.

The second trial started the spare satellites in DPO instead of LEO, to determine if this cislunar orbit might be viable for storage if the cost of moving the satellites there in the first place was not taken into account. The results of this trial are shown in Figure 20. Compared to the same trial conducted with only Earth-launched refueling missions, little has changed. The lower delta-V between DPO and LLO compared to DPO and LEO has allowed the service vehicle to move all needed satellites to LLO in order to use it as the sole parking orbit.

The third trial returned the satellites to LEO at the start, but moved the launchers to LEO and LLO. In this manner, this trial eliminated the fuel launch cost from consideration. The results of this are shown in Figure 21. Despite the elimination




Figure 19. Lunar ISRU Trial 1: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. It took 3.4 seconds to optimize.




Figure 20. Lunar ISRU Trial 2: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. Spare satellites started in DPO. It took 3.3 seconds to optimize.

of launch costs, of which the Earth launch had far more, the service vehicle in this network still prefers to use the lunar fuel depot. This is due to the markedly lower delta-V requirements from LLO to GEO than from LEO to GEO.

The fourth trial lowered the maximum event layer time to 3 days. The results of this are shown in Figure 22. Subsequently, the event layer limit was lowered to 1 day. The results of this are shown in Figure 23. From these trials, it can be deduced that in this scenario, it is cheaper to utilize the more fuel-intensive, time-optimized trajectories for lunar refueling than to bear the cost of an Earth launch. The only instance where the Earth-launched fuel becomes preferable is when the event layer is so short that transfers to LLO are impossible.

As a final trial for lunar refueling, a test was run with the fuel depot and satellites starting in DPO; this was to confirm that it was the mass of the depot/launcher, rather than the characteristics of LLO, that caused the optimizer to use its refueling orbit as a parking orbit. The results of this are shown in Figure 24. This trial demonstrated that, in this scenario, the optimal parking orbit is still the same as the orbit where the fuel depot starts or must travel to; even when that orbit is not LLO. Of interest is that moving the depot to DPO has resulted in transfers through L1 to save fuel.

The lunar ISRU trials make a very strong case for lunar ISRU as a factor that would make cislunar parking orbits practical for R3 of Earth-centric constellations. Given the same delta-V, TOF, and commodity data as the Earth-fuel-only trials, the network optimizer consistently chose the lunar fuel and brought all the satellites that it needed out to the LLO orbit of the fuel depot. This indicates that the combination of low launch costs and low transfer costs make lunar ISRU a game-changing factor that would make cislunar parking practical for R3 in this scenario.




Figure 21. Lunar ISRU Trial 3: this network solution was produced when the optimizer could refuel using fuel starting in LEO and LLO. It took 1.3 seconds to optimize.




Figure 22. Lunar ISRU Trial 4: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. The maximum event layer was shortened to 3 days. It took 4.5 seconds to optimize.




Figure 23. Lunar ISRU Trial 5: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. The maximum event layer was shortened to 1 day. It took 0.9 seconds to optimize.




Figure 24. Lunar ISRU Trial 6: this network solution was produced when the optimizer could refuel using fuel starting in DPO. Spare satellites started in DPO. It took 0.3 seconds to optimize.

4.3 L1 Lyapunov Constellation Included

In the third round of trials, this investigation expanded upon the original research questions, which considered only lunar ISRU and varied inclination as potential drivers for the use of cislunar orbits, with the inclusion of a hypothetical constellation in an L1 Lyapunov orbit in need of R3. In these trials, the demand in event layers 2 and 12 in Table 4 was shifted to the L1 node. These trials included both Earth and Moon-based refueling missions.

As in Sections 4.1 and 4.2, the first of these trials consisted of the default parameters with the L1 constellation added. The results of this are shown in Figure 25. This network demonstrated that with the inclusion of the L1 constellation, the use of lunar fuel is still cheaper than fuel from Earth, and that the optimal parking orbit for spare satellites is LLO.

The second trial with an L1 constellation mimicked the first, but with the lunar refueling mission removed. The results are shown in Figure 26. As with previous trials involving only Earth-launched refueling, the optimal solution for this scenario involves minimizing the fuel carried by refueling at every possible step, except at the end of the network when the service vehicle “knows” that it won’t need more fuel.

The third trial with an L1 constellation placed the spare satellites in DPO at the start of the network. Both Earth and Moon-based refueling missions were allowed. The results are shown in figure 27. As similar trial with only Earth-based missions allowed was also conducted; its results are in Figure 28. This trial showed that in this situation, DPO is largely not useful as a parking orbit. It is important to note that in both Figures 27 and 28, the service vehicle rarely actually visits DPO, preferring to pick up all its satellites at the beginning and move them to whichever orbit it will have to visit for refueling. Also of importance is that the DPO visit in event layer 13




Figure 25. L1 Constellation Trial 1: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. It took 3.4 seconds to optimize.




Figure 26. L1 Constellation Trial 2: this network solution was produced when the optimizer could only launch a refueling mission from Earth. A constellation has been simulated about the L1 Lagrange point. It took 4.4 seconds to optimize.

is simply to reduce delta-V; the commodities remaining there from event layer 12 are simply unused satellites (there is no demand for new satellites after step 12).

The final trial in this network involved both refueling missions enabled, with the event layer reduced to 3 days, the results are shown in figure 29. The most new information offered by this trial is that travelling through DPO in transfers around the Moon continues to be a viable way to save fuel in this scenario, even when the compressed schedule forces the use of time-optimized trajectories. Once again, the service vehicle moves all the satellites it needs to LLO in the first event layer and then picks them up as needed.

Overall, the addition of an L1 constellation in need to R3 does not appear to change the viability of various cislunar parking orbits. The network optimizer consistently chooses whatever orbit the fuel depot can travel to first as the parking orbit.

4.4 Miscellaneous Parameter Variations

To conclude this research, several trials were conducted to evaluate the effects of adjusting the optimizer parameters and the commodity properties on the network. All trials in this segment allowed both Earth and Moon-based refueling and included the L1 constellation. This enables the use of Figure 25 as a reference point, with its computation time of 4.1 seconds.

In order to identify potential vectors for optimizer performance improvement, several tests were conducted to determine the effects of the optimizer parameters on optimal network results and solve time. The number of possible parameters to vary is immense [16], but the two with the most relevance to the relatively simple network are MIPGap (Mixed Integer Programming Gap) and IntFeasTol (Integer Feasibility Tolerance), which control how close to a minimum or maximum a problem must be




Figure 27. L1 Constellation Trial 3: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. Spare satellites stored in DPO. It took 3.3 seconds to optimize.




Figure 28. L1 Constellation Trial 4: this network solution was produced when the optimizer could only launch a refueling mission from Earth. A constellation has been simulated about the L1 Lagrange point. Spare satellites stored in DPO. It took 10.7 seconds to optimize.




Figure 29. L1 Constellation Trial 5: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. The maximum event layer was shortened to 3 days. It took 1.8 seconds to optimize.

to be considered “optimized” and how close to an integer a continuous value must be during the branch and bound process.

The first of these trials consisted of increasing IntFeasTol from its default value of $1e-05$ to $1e-04$. The result of this is shown in figure 30. This result is not reflective of reality. Examination of the quantity of commodities thru each arc showed that the higher tolerance for non-integers allowed fuel to be stored at nodes that it could not normally occupy, by way of small parts of fuel depots remaining there. A second trial was conducted with the tolerance raised to $1e-03$, but despite a further reduction in computation time 10 1.7 seconds, the same error made the result unusable.

The second round of trials altered the MIPGap from its default value of $1e-04$. The results are shown in Figures 31, 32, and 33. Compared to a “control” optimization time of 3.0 seconds for a MIPGap of $1e-04$, these results indicate that adjusting the MIPGap tolerance can significantly affect the optimal solution with a relatively small change in optimization time. The results when the tolerance is raised to $1e-02$ are




Figure 30. Integer Feasibility Modification Trial 1: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. IntFeasTol was raised to 1e-04. It took 2.7 seconds to optimize.

questionable, and the results when raised to 1 are obviously unusable, but the results of tightening the tolerance to 1e-06 suggest that the DPO transfer may not be the best way for the service vehicle to save fuel in this instance.

Another example of the impact that the MIPGap tolerance can have on a solution is shown in Figures 34 and 35. This represented an attempt to replicate the trial in Figure 20, but with the spare satellites starting in LLO instead of DPO. This error specifically is the reason that the MIPGap tolerance was even tighter in the main trials; it demonstrates that fairly egregious errors can occur at the default value. In Figure 34, the service vehicle, with no satellites aboard, makes an unnecessary trip to L1 in event layer 4. There is no demand for any commodities at this node, and it is not saving fuel on a transfer, as it stops in LLO to refuel. This unnecessary trip appears consistently until the MIPGap solution tolerance is tightened in Figure 35.




Figure 31. MIPGap Modification Trial 1: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. MIPGap was raised to 1e-02. It took 2.5 seconds to optimize.




Figure 32. MIPGap Modification Trial 2: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. MIPGap was raised to 1. It took 1.0 seconds to optimize.




Figure 33. MIPGap Modification Trial 3: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. MIPGap was lowered to 1e-06. It took 3.2 seconds to optimize.




Figure 34. MIPGap Modification Trial 4: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. Spare satellites start in LLO. MIPGap was at its default value of 1e-04. It took 3.4 seconds to optimize.




Figure 35. MIPGap Modification Trial 5: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. Spare satellites start in LLO. MIPGap was lowered to $1e-06$. It took 2.5 seconds to optimize.

Trials were also conducted to evaluate the effects of varied satellite mass. In figure 36, the satellite mass has been increased to 1000 kg, in figure 37, it has been decreased to 10 kg.

The results of these trials demonstrate the same trend exhibited by the fuel launchers and the service vehicles; when a commodity is lower-mass, it becomes cheaper to carry it along and deposit it as needed. A heavier commodity becomes cheaper to leave in place until it is needed, with the service vehicle carrying only one at a time.

Overall, modifying the commodity and optimizer parameters is shown to have a significant effect on network performance. The optimizer parameters have the potential to shave off small amounts of computation time when tolerances are loosened, but can also lead to misleading results. Changing commodity mass and capacity can also affect the optimal result significantly; repeating all trials with different vehicle masses, for instance, would almost certainly produce different results; with a general tendency towards minimizing the movement of heavy commodities and “carrying along” lighter ones.




Figure 36. Satellite Mass Modification Trial 1: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. Satellite mass has been increased to 1000 kg. It took 6.2 seconds to optimize.




Figure 37. Satellite Mass Modification Trial 2: this network solution was produced when the optimizer could launch refueling missions from Earth and the Moon. A constellation has been simulated about the L1 Lagrange point. Satellite mass has been decreased to 10 kg. It took 3.0 seconds to optimize.

4.5 Final Analysis of Results

With all trials performed, the research questions have been answered, at least within the limitations of the delta-V/TOF “map” used for this research A. ED-GMCNF formulation is a valid and effective means of evaluating R3 architectures. Given the values for delta-V, TOF present in this research, cislunar orbits are generally not a practical solution for reducing inclination change costs for a multi-planar R3 regiment, but they are absolutely practical if lunar ISRU becomes a reality. The trial run with an increased satellite mass shown in Figure 36 suggests that some cislunar orbits could become practical staging orbits without the use of lunar fuel in different circumstances.

4.6 Chapter Summary

In this chapter, the methods discussed in Chapter III were expanded to cover a larger R3 network, and used to establish a definitive answer to the research questions.

V. Conclusions and Future Work

5.1 Conclusions

This investigation sought to combine the most recent capabilities of the ED-GMCNF network flow problem with the most recent advances in cislunar trajectory optimization in order to determine the practicality of cislunar orbits for R3 campaigns for Earth-orbiting satellite constellations.

After several trials, the research questions were answered, at least in the context of the provided delta-V and TOF map. An ED-GMCNF model is a suitable and effective way to model an R3 mission or series of missions. Cislunar orbits are not practical for servicing Earth-centric constellations under normal circumstances, but they become so when lunar ISRU is added to the mix. Additional trials revealed that the presence of a hypothetical cislunar constellation could also make these orbits practical staging points.

5.2 Limitations of the ED-GMCNF

Investigation beyond the toy problem revealed several limitations of the ED-GMCNF that had not been apparent until more complex networks were analyzed. These limitations color the conclusions of this investigation and form the basis of future work.

By requiring the delta-V and TOF of an arc to be determined ahead of time, the ED-GMCNF as presented in Ho [21] forces trade-offs between delta-V and TOF to be settled outside of the optimizer. The works of Brick and Ostman [3], [35] provided Pareto fronts for making these determinations, but incorporating the model functions of these fronts into the optimizer remains impractical in the current formulation; incorporating them directly would require a new arc to be created to correspond to

every data point, which would cause significant slowdown in networks with more than a handful of nodes.

Nodal motion, the actual movement of network nodes through space, is inherent in any space logistics network, but the effects become far more severe when the area of operations expands beyond Earth-centric orbits. When transfer times, and fuel costs vary over time, the use of constant values and functions to represent delta-V and TOF of arcs between them becomes less and less practical. This is, unfortunately, an area where the ED network is at a disadvantage compared to a TE network: accurately computing the transfer costs of an arc requires knowledge of when that arc starts, but this would be impossible for any arc except those starting in the first event layer, since the length of each event layer is not defined until travel occurs within it. In short, attempting to incorporate the orbital motion of the moon into the current ED-GMCNF would place the optimization function in a situation where it must know A to compute B, but must know B to compute A. In this situation, the optimizer only computes A. A potential solution to this problem, fixed point iteration, is discussed in section 5.3.

The time constraint by itself limits the network model in that it requires the user to have considerable knowledge of how a mission “should” play out. This was touched on in Chapter IV, where a decision to force the lower stage into a hold arc rather than the cargo was based on an assumption of how the mission would play out. In this investigation; all event layers were assigned the same length for simplicity. However, a real mission could very well have different maximum times depending on the function being performed. It could be very practical, for example, to allow for very long transfers out to a parking orbit to save fuel, but to require the fastest route possible to be used when a satellite has ceased functioning and needs a replacement. Doing this, of course, would require the operator to know which event

layers would include which arcs. This would not be a significant impediment on the relatively simple networks and missions that were considered in this investigation, but for a larger network, potentially servicing hundreds of nodes rather than tens, having to plot out what “should” happen before running to optimizer would not only represent a significant manpower cost for the user agency, but could potentially leave the optimizer blind to a more efficient solution than what the operator has in mind.

A further factor requiring foreknowledge of how missions must play out is the fact that all demands are permanent. As shown in both the test and main networks, commodities that satisfy demands disappear from the network. This is a convenient simplification when the commodities in question are replacement satellites and fuel; which mission planners would not want to be treated as an expendable resource once they were installed in a constellation. The disappearance of commodities post-demand becomes a problem, however, when repair needs to be considered. Requiring a service vehicle at a constellation to perform maintenance is an act that places demand on the service vehicle, rather than something it carries. This causes the service vehicle to disappear from the network and prevents future missions from occurring. It is possible to counteract this by supplying a replacement at the next event layer in the same node, but this is a crude solution that, as shown in Figure 38, leaves a gap in the network and also necessitates “carrying over” the service vehicle’s propellant and spare satellites in the same manner. The “spare parts” solution used in the trials of this investigation is less crude, but still not accurate to real life.

In the opposite vein of problems requiring foreknowledge of missions, the ED-GMCNF gives the service vehicle (and all other vehicles) a foreknowledge that would not exist in a real R3 scenario: when and where maintenance will be needed (and when it will not). Due to the mitigation mentioned in Chapter III, wherein a service vehicle is required to stay with a satellite in a constellation, this is not insurmountable. The




Figure 38. Network Solution with “Repair Gap”: this network solution demonstrates the “repair gap” (at the GEO node between event layers 10 and 11) caused when a demand is placed on a service vehicle. From the perspective of the network optimizer, the vehicle which moves the final satellite to the inclined GEO orbit is a brand new one; the previous one has been “used up”.

service vehicle must still travel the same amount that it would if it were responding to real, unexpected, R3 needs. However, this problem did consistently manifest in the final event layers of all trials. The service vehicle, “knowing” that the network would end and it would not have to refuel again, would choose a different path to its destination and would not carry enough fuel to return to the depot. The effects of this problem can also be seen in the unneeded spare satellites that are left in whatever orbit they start in, rather than carried to an appropriate parking orbit. Mitigation of this could be attempted by placing a demand in the parking orbit at the final event layer for excess commodities, but that would require the operator to know how many of each commodity are left, and where that parking orbit should be. However, this solution could preclude the instances where a parking orbit emerged “organically” (such as in Figure 36). Ultimately, this problem can be worked around by simulating a longer R3 campaign than needed and “cutting off” the end when deciding which results to use for mission planning.

Finally, there appeared to be an error in which the solved network would include the movement of partial satellites; i.e. satellites would be treated as a continuous value rather than a semi-integer value. This error only displayed itself once, when the first attempt was made to add a cislunar constellation to the network. This error disappeared when the fuel capacity was removed from the satellites (i.e. they could hold 0 kg of propellant rather than 20). The network still allowed the delivery of fuel to constellations (it would just disappear afterward), and it did not appear to affect results, so all trials were conducted again with this modification in place. This appears to be an error with the optimizer itself, since these solutions should not have been allowed. The exact cause is unknown and merits further investigation.

The final limitation of the current network formulation is a weakness that carries over from Ho [21]: the necessity for each event layer to be unidirectional to prevent flow generation loops. In the main network, commodities may only move right (to higher numbered nodes) in odd event layers, and may only move left (to lower numbered nodes) in even event layers. The direction and node numbering are arbitrary, but this “zig-zag” is necessary to prevent situations like the one shown in Figure 39. Note that this is from an earlier version of the final network, hence the smaller number of nodes and event layers. In these situations, since all movement arcs are computed at once; a service vehicle with no fuel is able to move as long as it passes through the fuel depot (LLO) at some point. This is called a flow generation loop, a term coined by Ho [21].

Overall, while the ED-GMCNF is capable of modelling an R3 mission campaign, there are several drawbacks that can be overcome to further improve its fidelity and utility for mission planners.




Figure 39. Flow Generation Loop Error: this network solution demonstrates the flow generation loop phenomena (in event layer 2) caused when unlimited travel is allowed with a event layer, regardless of direction. From the perspective of the network optimizer, the vehicle needs a total amount of fuel to complete the event layer, and it does not matter if it actually has that at the start, as long as it can acquire it. The service vehicle starts with no fuel but completes the network anyways because it travels to and from the LLO fuel depot in a single step. [Graphic generated in MATLAB]

5.3 Future Work

Proposals for future work in this area of research are based on the difficulties outlined in the previous section.

The error which was causing satellites to act as a continuous commodity rather than a semi-integer commodity did not affect the operation or utility of the ED-GMCNF in any major way once the satellite fuel capacity was removed, but the cause should still be identified.

A possible solution for the restrictions imposed by the time constraints is to automatically restrict every commodity to moving, at most, over one movement arc per event layer. This may increase computation time, but it would eliminate the need for a fair amount of planning ahead. This would solve the problem of the event layers having to be uni-directional to prevent flow generation loops.

The problems caused by the network having foreknowledge of what commodities will be required when is best solved by simply not providing more commodities than the trial will need, as shown in Figures 40 and 41. Unfortunately, this is not a practical solution for continuous commodities such as propellant, which is still dropped to lighten vehicles on a routine basis. For instance, in Figure 40, the service vehicle dumps almost 3500 kg of propellant at the beginning of the network before departing LEO. This is a difficult problem to solve because the “dumped” commodities do not travel along an arc; therefore they incur no cost. Ultimately, the most practical solution to this problem appears to be estimating initial amounts of required fuel; the problem becomes much less severe after the service vehicle’s first journey to the depot. At this point, the vehicle automatically picks up the amount of propellant it needs.

Physical motion of network nodes, such as that of the Moon through its orbit around Earth, remains the most difficult of the aforementioned problems with the ED-GMCNF formulation to solve, since unlike the low-thrust curves and time constraints, which have brute force solutions within the existing constraints of the network, enabling nodal motion does not. Mathematics suggests that in a situation amounting to a complex system of simultaneous equations, the best solution would be to pursue an iterative solution. A possible solution is to restrict movement to single arc per event layer and use fixed-point iteration techniques to update values of delta-V and TOF after predefined time increments have passed.. One potential method is the use of dynamic programming. More research is needed to determine the applicability of this to this particular problem.

In addition to addressing problems, the incorporation of low-thrust trajectories into the network is a logical part of any next step, as mentioned in Chapter II.




Figure 40. Extra Satellites Removed Example 1: this network is a replica of Figure 29, but with only 5 satellites provided at the start instead of 10. [Graphic generated in MATLAB]




Figure 41. Extra Satellites Removed Example 2: this network is a replica of Figure 28, but with only 5 satellites provided at the start instead of 10. [Graphic generated in MATLAB]

5.4 Contributions of this Research

This research represents the first formulation of a space-centric ED-GMCNF with the capacity to select from fuel and time-optimized arcs between the same node pair, and is a significant reduction in independent variables from the previous work of Ho [21]. It is also the first of these networks to model a reusable launch vehicle.

5.5 Summary

This chapter summarized the conclusions of this investigation, the weaknesses of the ED-GMCNF network as a tool and potential work that could be done to mitigate those weaknesses. The Ed-GMCNF proved an effective tool for mission analysis, and demonstrated, at least within the scope of the assumptions and provided data, that cislunar orbits can effectively play a roll in satellite constellation R3 under the correct circumstances, namely the use of lunar ISRU and when satellite mass is increased tenfold. This research represents the first formulation of a space-centric ED-GMCNF with the capacity to select from fuel and time-optimized arcs between the same node pair, and is a significant reduction in independent variables from the previous work of Ho [21]. It is also the first of these networks to model a reusable launch vehicle.

Appendix A. Main Network delta-V and TOF Table

Node 1	Node 2	Type	ΔV (km/s)	TOF (days)
L1	L2	Min Fuel	0.000848	26.02
L1	DPO	Min Fuel	0.042	5.2213
L1	DPO	Min Time	0.5094	0.6465
L1	GEO	Min Fuel	1.4128	4.3012
L1	GEO	Min Time	6.0123	0.9273
L2	GEO	Min Fuel	1.6995	5.1525
L2	GEO	Min Time	5.9896	1.4188
LEO	GEO	Min Fuel	4.5084	0.2139
LEO	GEO	Min Time	5.5812	0.2075
LEO	PGEO	Min Fuel	5.6403	0.2191
LEO	PGEO	Min Time	6.4111	0.2139
L1	LEO	Min Fuel	4.0263	4.2593
L1	LEO	Min Time	5.838	1.3272
L2	LEO	Min Fuel	4.0417	6.0879
L2	LEO	Min Time	5.8718	1.9059
L1	PGEO	Min Fuel	2.7763	4.9236
L1	PGEO	Min Time	3.4683	2.2708
L2	PGEO	Min Fuel	2.8767	7.0194
L2	PGEO	Min Time	3.7405	2.6576
L2	DPO	Min Fuel	0.1825	5.3699
L2	DPO	Min Time	0.4632	0.9664
DPO	LEO	Min Fuel	3.6456	5.3874
DPO	LEO	Min Time	5.3286	1.7268
DPO	GEO	Min Fuel	1.5525	7.9344
DPO	GEO	Min Time	3.5355	1.8349
DPO	PGEO	Min Fuel	1.8066	7.8211
DPO	PGEO	Min Time	3.9151	1.70535
DPO	LLO	Min Fuel	0.94	1.097
DPO	LLO	Min Time	1.7998	0.312
LLO	LEO	Min Fuel	3.9332	4.4656
LLO	LEO	Min Time	5.4587	1.4415
LLO	GEO	Min Fuel	1.7856	5.8645
LLO	GEO	Min Time	3.0613	1.8646
LLO	PGEO	Min Fuel	2.6895	5.6612
LLO	PGEO	Min Time	3.6273	1.66309
L2	LLO	Min Fuel	0.7517	6.6692
L2	LLO	Min Time	0.8607	4.3279
L1	LLO	Min Fuel	0.8472	2.0027
L1	LLO	Min Time	0.9324	1.4337
Earth	LEO	Launch	9.4	0.02
Moon	LLO	Launch	1.73	0.02
GEO	PGEO	Min Fuel	3.9272	0
PGEO	GEO	Min Fuel	3.9272	0

This information was used in all main network trials. Note that all arcs were assumed to be reversible.

Bibliography

- [1] *Advanced Aerospace Medicine On-line*. Federal Aviation Administration. 2018.
- [2] Jonathan Amos. *Space Harpoon Skewers Orbital Debris*. BBC. Feb. 2019. URL: <https://www.bbc.com/news/science-environment-47252304>.
- [3] John N. Brick. “Military Space Mission Design and Analysis in a Multi-Body Environment: An Investigation of High-Altitude Orbits as Alternative Transfer Paths, Parking Orbits for Reconstitution, and Unconventional Mission Orbits”. MA thesis. Air Force Institute of Technology, 2017.
- [4] John Carrico et al. “Launch Opportunity Analysis of GEO Transfer with High Inclination using Lunar Gravity Assist”. In: Aug. 2019.
- [5] Jean-Francois Castet and Joseph H. Saleh. “Satellite and satellite subsystems reliability: Statistical data analysis and modeling”. In: *Reliability Engineering System Safety* 94.11 (2009), pp. 1718–1728. DOI: <https://doi.org/10.1016/j.ress.2009.05.004>. URL: <http://www.sciencedirect.com/science/article/pii/S0951832009001094>.
- [6] *Challenges to Security in Space*. Defense Intelligence Agency. 2019.
- [7] Y. L. Chen and Y. H. Chin. “Multicommodity Network Flows with Safety Considerations”. In: (July 1990). DOI: <https://doi.org/10.1287/opre.40.1.S48>.
- [8] Stephen Clark. *NASA: Tracking CubeSats is Easy, but Many Stay in Orbit Too Long*. Spaceflight Now, 2015. URL: <https://spaceflightnow.com/2015/07/30/nasa-tracking-cubesats-is-easy-but-many-stay-in-orbit-too-long/>.
- [9] *Competing In Space*. National Air and Space Intelligence Center. 2018.
- [10] John Connolly. *Human Lunar Exploration Architectures*. NASA. Oct. 2012.
- [11] Ian A. Crawford. “Lunar Resources: A Review”. In: *Progress in Physical Geography* (2014).
- [12] Jacob A. Dahlke. “Optimal Trajectory Generation in a Dynamic Multi-Body Environment Using a Pseudospectral Method”. MA thesis. Air Force Institute of Technology, 2018.
- [13] Gregory F. Dubos, Jean-Francois Castet, and Joseph H. Saleh. “Statistical reliability analysis of satellites by mass category: Does spacecraft size matter?” In: *Acta Astronautica* 67.5 (2010), pp. 584–595. DOI: <https://doi.org/10.1016/j.actaastro.2010.04.017>. URL: <http://www.sciencedirect.com/science/article/pii/S0094576510001347>.
- [14] Tom Fleischman. *Cornell’s Quest: Make the First CubeSat to Orbit the Moon*. Cornell Chronicle, 2016. URL: <https://news.cornell.edu/stories/2016/09/cornells-quest-make-first-cubesat-orbit-moon>.

- [15] Philip E. Gill, Walter Murray, and Michael A. Saunders. *User's Guide for SNOPT Version 7: Software for Large-Scale Nonlinear Programming*. Stanford University and UC San Diego, 2018.
- [16] *Gurobi Documentation*. Gurobi Optimization. Oct. 2019. URL: <https://www.gurobi.com/documentation/8.1/refman/parameters.html>.
- [17] Koki Ho. “Dynamic Modeling and Optimization for Space Logistics Using Time-Expanded Networks”. In: *Acta Astronautica* 105 (2014), pp. 428–443.
- [18] Takuto Ishimatsu. “Generalized Multi-Commodity Network Flows: Case Studies in Space Logistics and Complex Infrastructure Systems”. PhD thesis. Massachusetts Institute of Technology, June 2013.
- [19] Takuto Ishimatsu et al. “Generalized Multicommodity Network Flow Model for the Earth–Moon–Mars Logistics System”. In: *Journal of Spacecraft and Rockets* 53.1 (2016), pp. 25–38. DOI: 10.2514/1.A33235. URL: <https://doi.org/10.2514/1.A33235>.
- [20] Alon Itai. “Two-Commodity Flow”. In: *Journal of Associated Computing* (July 1978), pp. 596–611.
- [21] Bindu Jagannatha and Koki Ho. *Event-Driven Space Logistics Network Optimization with Embedded Propulsion Technology Trades*. 2019.
- [22] Bindu Jagannatha and Koki Ho. “Event-Driven Space Logistics Network Optimization with Embedded Propulsion Technology Trades”. In: *ArXiv* abs/1904.09364 (2019).
- [23] Pauline Jakob et al. “Optimal Satellite Constellation Spare Strategy Using Multi-Echelon Inventory Control”. In: *Journal of Spacecraft and Rockets* 56.5 (2019), pp. 1449–1461.
- [24] John P. Mayberry Joshua P. Davis and Jay P. Penn. *On-Orbit Servicing: Inspection, Repair, Refuel, Upgrade, and Assembly of Satellites in Space*. The Aerospace Corporation. 2019.
- [25] Warren B. Powell Judith M. Farvolden and Irvin J. Lustig. “A Primal Partitioning Solution for the Arc-Chain Formulation of a Multicommodity Network Flow Problem”. In: *Operations Research* 41 (4 Aug. 1993), pp. 622–806.
- [26] Wang Sang Koon et al. *Dynamical Systems, the Three-Body Problem and Space Mission Design*. 1.2. California Institute of Technology, Apr. 2011.
- [27] David Kornuta et al. “Commercial lunar propellant architecture: A collaborative study of lunar propellant production”. In: *REACH* 13 (2019), p. 100026. ISSN: 2352-3093. DOI: <https://doi.org/10.1016/j.reach.2019.100026>. URL: <http://www.sciencedirect.com/science/article/pii/S2352309318300099>.

- [28] Ailsa H. Land and Alison G. Doig. “An Automatic Method for Solving Discrete Programming Problems”. In: *50 Years of Integer Programming 1958-2008: From the Early Years to the State-of-the-Art*. Ed. by Michael Jünger et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 105–132. ISBN: 978-3-540-68279-0. DOI: 10.1007/978-3-540-68279-0_5. URL: https://doi.org/10.1007/978-3-540-68279-0_5.
- [29] Edward K. Morlok Larry J. LeBlanc and William P. Pierskalla. “An Efficient Approach to Solving the Road Network Equilibrium Traffic Assignment Problem”. In: *Transportation Research* (Oct. 1975).
- [30] Juseong Lee et al. *Pareto front generation with knee-point based pruning for mixed discrete multi-objective optimization*. 2017. arXiv: 1709.09311 [math.OC].
- [31] Andrew S. LeValley. “A Mixed Integer Programming Framework for the Fuel Optimal Guidance of Complex Spacecraft Rendezvous and Proximity Operations Missions”. MA thesis. Air Force Institute of Technology, 2019.
- [32] Federico Malucelli. *Integer Linear Programming*. Politecnico di Milano. URL: <http://home.deib.polimi.it/malucelli/didattica/appunti/4ENG.pdf>.
- [33] James Mattis. *Summary of the 2018 National Defense Strategy*. 2018.
- [34] Murat Oğuz, Tolga Bektaş, and Julia A. Bennell. “Multicommodity flows and Benders decomposition for restricted continuous location problems”. In: *European Journal of Operational Research* 266.3 (2018), pp. 851–863. ISSN: 0377-2217. DOI: <https://doi.org/10.1016/j.ejor.2017.11.033>. URL: <http://www.sciencedirect.com/science/article/pii/S0377221717310445>.
- [35] Joshua A. Ostman. “Cislunar Trajectory Generation with Sun-Exclusion Zone Constraints Using a Genetic Algorithm and Direct Method Hybridization”. MA thesis. Air Force Institute of Technology, 2019.
- [36] Russell E.A. Smith and J. W. Yound. “Trajectory optimization for an Apollo-type vehicle under entry conditions encountered during lunar return”. In: 1967.
- [37] *Spacecraft Trajectory Optimization*. Cambridge Aerospace Series. Cambridge University Press, 2010. DOI: 10.1017/CB09780511778025.
- [38] Victor Szebehely. “Theory of Orbit”. In: Academic Press, 1967. ISBN: 978-0-12-395732-0. DOI: <https://doi.org/10.1016/B978-0-12-395732-0.50007-6>. URL: <http://www.sciencedirect.com/science/article/pii/B9780123957320500076>.
- [39] Andre Tartar and Yue Qiu. *The New Rockets Racing to Make Space Affordable*. Bloomberg Businessweek. July 2018. URL: <https://www.bloomberg.com/graphics/2018-rocket-cost/>.
- [40] C. Taylor et al. *A Mathematical Model for Interplanetary Logistics: Logistics Spectrum*. English. Other. Type: Magazine. 2007.

- [41] I-Lin Wang. “Multicommodity Network Flows: A Survey, Part I: Applications and Formulations”. In: 15 (Dec. 2018), pp. 145–153. DOI: 10 . 6886 / IJOR . 201812_15(4) . 0001.
- [42] Ryan J. Whitley et al. “EARTH-MOON NEAR RECTILINEAR HALO AND BUTTERFLY ORBITS FOR LUNAR SURFACE EXPLORATION”. In: 2018.
- [43] William E. Wiesel. *Modern Orbit Determination*. Beavercreek, Ohio: Aphelion Press, 2010. ISBN: 978-1453611982.
- [44] William E. Wiesel. *Spaceflight Dynamics: Third Edition*. Aphelion Press, 2010.
- [45] Richard D. Wollmer. *Multicommodity Supply and Transportation Networks with Resource Constraints: The Generalized Multicommodity Flow Problem*. RAND Corporation. Mar. 1970.